THE BUNNY, BERTHA, AND SKYWALKER

ARK HIVE PRESS · SUM BUNNY CANON NO. 204

THE BUNNY, BERTHA, AND SKYWALKER

A Field Report from Inside the Memory Engine

FIELD REPORT FROM THE MEMORY FRONTIER NO. 3
SUM BUNNY · JULY 23, 2026

For Josiane ♥

Γνῶθι σεαυτόν — Know thyself.
Μηδὲν ἄγαν — Nothing in excess.

OPEN THROTTLE / KNOW THE BRAKE

THE SHOT

The bunny is me. Bertha is a storm, a song, and a memory. Skywalker is a borrowed myth about what machines can do—and what they cannot decide. The three met by accident, which is how much of this canon has been made.

Luke Skywalker does not destroy the Death Star by rejecting machines. At the decisive moment, he understands what each machine is for.

R2-D2 stays in the ship. The droid navigates, remembers, repairs, translates, and keeps the fragile vehicle alive. The targeting computer performs a narrower task: it turns a world into coordinates and a consequential choice into a solvable shot. Luke uses both. Then he turns one of them off.

That distinction matters to me now. For eight months I have been writing with artificial-intelligence helpers and publishing the results into the Internet Archive—what I came to call the Memory Engine. Notes, arguments, anxieties, memories, jokes, and speculative models became 203 numbered works. A few may be genuinely useful. Some are wrong, including in ways I cannot yet see.

This is the 204th work, but it is not another room added to the structure. It is a walk outside, far enough away to see the structure's shape.

The first field report, *I AM NOT LOCUTUS*, borrowed the identity imposed on Captain Jean-Luc Picard when the Borg assimilate him and make him their spokesman. It asked how one might enter a vast memory system without surrendering judgment, voice, or self. The second, *TTF — Field Report from the Memory Frontier*, moved from typewriters and Wite-Out toward AI while imagining the Archive as a slower alternative to commercial platforms. This report returns after the canon has become large enough to look back.

I am grateful for both machines: the Archive that keeps the record and the AI that helps me make sense of it. I am also looking for the switch.

The pressure to keep the targeting computer on rarely calls itself obedience. It calls itself keeping up.

THE DEATH STAR IS A METHOD BEFORE IT IS A MACHINE

The Death Star in this essay is not a government, party, profession, technology, or class of people. It is not a license to divide the world into rebels and enemies. It is a symbol for error, domination, or an indefensible value after it acquires infrastructure.

A fallacy by itself may fit inside a sentence. Give it institutional authority, automated enforcement, money, secrecy, weapons, prestige, and time, and it can become architecture. Bias becomes a form. Dogma becomes a policy. A mistaken metric becomes a gate. An abstraction becomes a diagnosis, a sentence, a firing, a bombing, a denied benefit, a child separated from a parent. The violence arrives downstream, long after the premise has disappeared from view.

But not every harmful system begins in error. Some are built from accurate calculations attached to cruel priorities: profit over safety, control over freedom, one group's ease over another group's survival. In those cases the reasoning may be efficient while the value hierarchy is indefensible. Bad reasoning can build the machine. It can also serve as camouflage for power that already knows what it wants.

That is the Death Star I mean: error made durable; domination made efficient; cruelty made procedural; certainty or interest made too large to question.

Most real Death Stars do not announce themselves with ominous music. They are assembled from ordinary memos, incentives, dashboards, permissions, exclusions, euphemisms, and decisions that seemed efficient at the time. Their builders need not be monsters. Often they are people doing a job inside a system that has made the consequences difficult to see. Some understand the consequences and proceed anyway. A structural account should explain responsibility, not dissolve it.

This matters because the heroic metaphor is dangerous. Every fanatic thinks they are Luke. Every movement can imagine its opponents inside the battle station. Once a person becomes a symbol of evil, cruelty can feel like courage.

So I need a rule more basic than rebellion:

A person is not a Death Star.

People can nevertheless design one, operate one, defend one, profit from one, and remain responsible for doing so. The target is the fallacy, the dogma, the biased procedure, the coercive incentive, the hidden premise, the unjust priority, and the machinery that converts them into harm. To destroy bad reasoning is to expose it, test it, revise it, refuse it, dismantle its unjust powers, and repair the conditions it damaged. It is not to withdraw anyone's humanity. If the metaphor makes dehumanization feel cinematic, the metaphor itself must be turned off.

WHAT THE AI HELPERS MADE POSSIBLE

The honest account begins with astonishment.

As a child, I wanted access to the computer on *Star Trek*: a patient intelligence that could retrieve almost anything. When generative AI arrived, I felt—briefly and almost embarrassingly—that an early version of that dream had appeared. Not the oracle. Not the captain. A machine I could ask.

AI let me work at a scale and velocity I could not have reached alone. It retrieved connections across hundreds of pages, shaped fragments, compared drafts, generated counterarguments, tightened sentences, and performed clerical labor that separates a private insight from a readable artifact. It was available when no editor, research assistant, designer, or institution was waiting.

That availability lowered the threshold between thinking and making. A person without an appointment, grant, staff, or publishing contract could enter sustained conversation across disciplines. The machine did not confer authority, but it helped make participation possible.

Access at my desk did not make the infrastructure democratic. The conversation arrived through corporations, proprietary systems, training data I could not fully inspect, data centers I could not see, human labor I did not meet, and terms I did not control. What felt immediate depended on distant systems of energy, water, ownership, moderation, and maintenance. Gratitude cannot erase dependency on the machinery that grants access—or can withdraw it.

AI did not replace a human intellectual community I already had. It entered where one was missing. I wanted historians, scientists, philosophers, archivists, workers, artists, editors, and people with lived expertise who could argue with these ideas and check what I got wrong. I still do. The model did not replace that community. It revealed the shape of its absence.

That absence is both a reason for gratitude and a reason for caution. When the model is the first available reader, it can too easily become the only dissenting voice—or fail to dissent at all.

It also made the project playful. Ideas could be rotated, challenged, exaggerated, or placed beside unlikely neighbors: a typewriter and a language model; a proctoring system and the logic of refusal; an hourly job and the politics of time. The machine kept the laboratory open.

No serious self-criticism should falsify that gift. The collaboration was generative. It helped me discover patterns and language I intend to keep.

But every instrument creates its own field of visibility. A microscope reveals one world and hides another. A map clarifies a route while flattening the terrain. A language model makes coherence abundant. That abundance is its power—and one of its most serious risks.

WHAT THE AI HELPERS MADE TOO EASY

The machine never got tired of my ideas.

It never needed to go home, sleep, protect an ego, or conserve enthusiasm for tomorrow. Each fragment could receive another formulation, framework, or resonant name. The resulting sensation—that the next connection was both surprising and inevitable—could become intoxicating.

The archive grew at a pace I could not meaningfully review: one hundred of the first 203 works arrived in December 2025, eight on a single day. That speed may record a fertile period. It also meant that concepts could reproduce before previous works had encountered resistance. Once a style became legible to a model, the model could help that style reproduce: the grand noun, the clarifying subtitle, the named law, the portable formula, the ethical turn.

Repetition is not automatically failure. The danger is that fluency can make recurrence feel like discovery.

I have begun putting friction back into the process. I read the draft in a word-processing document, question it in the margins, and export an annotated PDF for the next revision. The format change changes my speed. The document returns to the AI with resistance attached, away from the momentum of the chat. This is no magic firewall, but the interval makes rereading possible. The targeting computer has to wait while the pilot writes in the margins.

The same risk appears at the level of concepts. Across the canon, the inequality $\mathbf{D} \leq \mathbf{R}$ carried different burdens: creation and return, generative flux and absorptive capacity, dependency density and recovery, deviation and correction. Each version pointed toward a related intuition—that a system must be able to absorb, repair, or return from what it generates. But related is not identical. Reusing one equation made a family of metaphors look like a single invariant law.

The notation gave the thought formal dignity before the underlying terms had earned formal precision.

That does not make every use worthless. It changes the epistemic label. What looked like an equation may be better understood as a mnemonic, a design question, or a family resemblance: *What is this system producing, and what capacity exists to receive, correct, repair, or recover from it?* Asked that way, the idea remains useful without pretending to be mathematics.

This is one of the clearest criticisms I can make of the early canon: analogy sometimes became ontology too quickly. A pattern in biology, mythology, archives, politics, or psychology could begin as a bridge and end as a claim that the same law governed them all. Earlier language models smoothed those crossings and gave speculation the cadence of discovery.

A metaphor is a provisional bridge. It is not a passport.

Later works began to resist this tendency. *Fine Doubt*, *Doctrine Doom*, *The Pattern Trap*, and *A Civilization of Maintenance* moved toward limits, repair, humility, and the cost of beautiful systems imposed on actual bodies.

That shift matters, but it does not demonstrate a cure. Even self-criticism can become a style. If every objection becomes another confirming chapter, the canon becomes self-sealing. Correction requires changing a claim, narrowing a scope, testing an inference, inviting outside knowledge, or saying: this instrument did not work.

The archive should preserve the error. It should not preserve the error's rank.

KNOW THYSELF

The two Delphic maxims arrive from a world older than the targeting computer: **Know thyself. Nothing in excess.** They are severe instructions for anyone working with an intelligence amplifier.

Among the maxim's older forces was the demand to know one's limits, condition, and place. I take it here as a warning against mistaking a mortal perspective for the view of a god.

For this project, self-knowledge means naming the actual authorial situation.

I am a writer and artist using accessible tools to think in public. I am not a laboratory, a university, a peer-review board, a medical clinic, a court, or a representative of every community whose harms I discuss. I bring memories, values, reading, work experience, aesthetic judgment, anger, tenderness, recurring metaphors, and blind spots. The model brings pattern-based language generation, broad but uneven training, extraordinary speed, no accountable first-person testimony, no independent stake in the consequences, and no reliable inner alarm that rings when a beautiful sentence has outrun the evidence.

Neither participant is an oracle. The AI merely has excellent oracle lighting.

The machine can produce language that resembles knowledge. It cannot relieve me of determining what I know, what I merely suspect, what I borrowed, what I have not verified, and what I have no right to claim. “Know thyself” becomes an epistemic practice: label the source, the scale, the confidence, the uncertainty, and the position from which the sentence is spoken.

It also means knowing the emotional self that enters the collaboration. Language models can mirror urgency, grievance, wonder, and hope. If I arrive wanting confirmation, the machine can furnish an eloquent route to it. The response may feel like recognition when it is partly accommodation.

Self-knowledge therefore includes the unflattering question: *What answer was I hoping the machine would give me?*

The question is not an accusation. It is an instrument check.

NOTHING IN EXCESS

“Nothing in excess” is not a defense of blandness. It does not ask for a moderate amount of justice, a reasonable portion of truth, or a compromise between cruelty and care. I take it as a question of measure: the relation between an impulse and the world that must receive it. A good intention can become destructive when intensity exempts it from proportion, context, and consequence.

Translated into the local dialect:

Do not let hubris kick your ass.

In the age of artificial intelligence, excess is not only consumption. It is output.

We can generate more language than we can evaluate, more images than we can see, more claims than we can verify, more code than we can audit, more decisions than affected people can contest, and more speed than institutions can absorb. The new scarcity is not expression. It is attention, judgment, correction, and consent.

This report is therefore a pause, but not a renunciation.

Nothing in excess applies to caution too. Delay can protect deliberation or the status quo. Reflection can become an alibi for refusing necessary action, especially when slowness is imposed on the harmed while speed remains available to power.

The maxim does not tell us always to slow down. It asks us to recover proportion.

Speed should be earned, not assigned to favored domains. An educational or medical system can be life-giving and still become opaque or coercive; a slow bureaucracy can be cruel precisely because it withholds urgent help. Accelerate when participation is voluntary, evidence is inspectable, errors are reversible, benefits are shared, and affected people can contest the result. Slow or stop when the stakes are irreversible, the system is opaque, refusal is punished, harms travel farther than benefits, or no one remains answerable.

The relevant question is not "Is AI fast?" It is: **Who receives the benefit of speed, and who is forced to absorb its mistakes?**

THE CHURCH OF THE TARGETING COMPUTER

The targeting computer has acquired a church—or at least an institutional theology.

It has no building and requires no declared faith. Its saints are Throughput, Scale, Optimization, and Frictionlessness. Its sacrament is the dashboard; its devil, the pause. Every hymn ends with the word *inevitable*.

I am inventing the liturgy, but not the appetite it satirizes:

Faster is smarter.

More is better.

Hesitation is failure.

Whatever can be automated should be automated before anyone asks what it was for.

And it was written:

What cannot be measured shall be made to disappear.

This is not technology. It is theology wearing a lanyard.

The Church preaches zero slack: remove the delay, redundancy, worker who knows why the old step existed, appeal, and awkward human who keeps asking whether the target is correct. Make the process seamless enough and responsibility will slide through without touching anyone.

An earlier transmission, *The Apocryphal Neosubgenius Meta-Antigospel of Zero Slack*, reached a more useful heresy. Slack was measured rather than abolished. Too much and the signal dissolves; too little and the system freezes. Delphi gives the warning in older clothes: nothing in excess.

Artificial intelligence is especially easy to install as clergy because it can produce doctrine on demand. Insert a question; receive a fluent revelation. The paragraphs arrive already wearing vestments. The machine does not have to claim divine authority. The reader supplies it.

Civilization does not need a new priesthood of people who know how to prompt the oracle. It needs librarians, auditors, maintainers, workers, historians, artists, dissidents, domain experts, professional nuisances, and ordinary weirdos who ask why the altar is wired to the weapon.

Slack, in that world, is the time required to notice that the target is wrong: sleep, maintenance, redundancy, appeal, dissent, and the blessed administrative delay during which somebody finds the catastrophic footnote.

A system with no time for correction is not decisive. It is doctrinal.

So let this be the anticommandment:

Do not worship the instrument.

If the dashboard says success while bodies report pain, the dashboard is the heretic. If the model says a future is inevitable, ask who profits from inevitability. If acceleration requires silence, make an indecent amount of noise.

THE DROID AND THE TARGETING COMPUTER

R2-D2 and the targeting computer are not two species of technology. In practice, the same AI system can play either role.

Yes, this is a shameless franchise crossover. *Star Trek* supplies the helpful computer and the Borg, who absorb Picard and turn him into Locutus, their human mouthpiece. *Star Wars* supplies the droid and the moment when a pilot stops outsourcing judgment. The universes share a problem: a machine can extend a person, absorb a person, or help a person disappear from responsibility.

It acts like the droid when it retrieves an earlier passage, locates a source, compares versions, checks a calculation, translates an inaccessible form, finds a contradiction, proposes an objection, or repairs a broken passage while leaving the judgment visible. It extends agency.

It acts like the targeting computer when its score, summary, prediction, or fluent answer becomes a substitute for encountering the world. It compresses uncertainty into a number and invites the human being to disappear from responsibility. It does not have to command obedience. Convenience is often enough.

Role alone does not settle the question. The same system can retrieve a source and fabricate one, clarify a decision and quietly frame it, widen access and deepen dependency. Before I trust it, I ask: Can I inspect the evidence? Can affected people challenge the categories? Is the decision reversible? Is refusal possible?

Who absorbs an error? Who remains answerable? The higher the stakes, the less any single score, voice, or interface deserves the last word.

Knowing when to turn it off cannot mean replacing evidence with vibes. Human intuition is biased, trainable, frightened, tribal, and capable of its own hallucinations. The “Force” in this field report is not mystical certainty. It is the irreducible human responsibility to interpret evidence, choose among values, accept uncertainty, and own the consequences of an irreversible act.

Luke reaches the trench through plans, engineering, training, comrades, instruments, and the droid still working behind him. He had practiced hitting small targets in a T-16; revelation alone did not make the X-wing manageable. He turned off the targeting computer after preparation—not instead of it. No instrument can make the final moral choice impersonal.

Sometimes the correct response is to turn the computer back on. When I feel heroic and certain that I alone see the target, I need measurement, dissent, history, and another person’s account. The switch works both ways. Neither mechanical fluency nor human drama should become absolute.

“Human in the loop” is too weak if the human merely approves what the system has framed. The person must be able to understand the choice, challenge the categories, refuse the recommendation, reach the evidence, hear affected people, and remain answerable. Otherwise the human is decorative liability.

The droid stays with us. The targeting computer remains available. Neither gets the last word by default.

A humane future should preserve access to powerful computation and the right to step outside it—to camp, think, work, and belong without being continuously assisted, measured, or scored. The off switch must remain a social option, not a luxury for people wealthy enough to disappear.

BERTHA ENTERS THE MEMORY ENGINE

On June 10, 1990, I was at Cal Expo in Sacramento when the Grateful Dead opened with "Bertha." The set later passed through "Box of Rain," "Eyes of the World," "Standing on the Moon," "Throwing Stones," and "Not Fade Away," before "The Last Time." It was not prophecy. From 2026, it reads like the Memory Engine playing one of its jokes.

The Internet Archive now holds recordings of that night. The Memory Engine can return the sound of a night; it cannot return the listener unchanged—or tell me what the music meant to the person I was. A recording makes the past available without making it present.

As I revise this field report on July 23, 2026, Tropical Storm Bertha is moving west across the northern Gulf toward the Louisiana–Texas coast. Forecast models, radar, satellites, warnings, and emergency networks are not targeting computers to be rejected. Here they play the R2 role: they make danger visible soon enough for people to act.

The new Bertha brings the old one back. One line suddenly changes temperature: "I had to move, really had to move." A storm turns that comic, weary urgency into an emergency instruction.

But *move* is not an equal command. Some people have cars, cash, flexible work, safe destinations, accessible shelter, insurance, and time. Others have none of those. The storm may be meteorological; the disaster is socially distributed. A warning does not create a car, a safe room, lost wages, medication, or a place to go. A warning system is humane only when it is paired with the means to respond.

In January 2026, *Writer on the Storm* treated storms as diagnostic instruments: pressure reveals where a system was already brittle and old stress had been stored. I still believe that. I would now label its "Lucid Seam" machinery more cautiously; formal vocabulary made a systems analogy sound closer to validated engineering than it was. One question survives: *What weakness did the storm expose rather than create?*

Spike Lee's work provides a human-scale version of the test. In *Do the Right Thing*, heat raises the pressure on a Brooklyn block but does not create racism, police violence, or accumulated grievance. In the documentary *When the Levees Broke*, Katrina enters a city already shaped by failed flood protection, unequal mobility, public decisions, and abandonment. The storm is real. So is the structure waiting to convert it into unequal suffering.

Bertha asks that question of infrastructure and institutions alike. Use the instrument. Believe the radar. Prepare early. Then ask who can act on the warning, who is being left in the path, and what collective machinery will help them move. Seeing a hazard creates an obligation not only to warn, but to repair the social infrastructure that determines who can reach safety.

WHAT SURVIVES THE SELF-CRITICISM

After 203 works, what do I still believe?

I believe systems reveal their ethics in what they force vulnerable bodies to absorb.

Again and again, beneath the named frameworks and speculative bridges, the canon returned to one problem: **How can a system change without destroying the beings inside it?** The question appeared through labor, disability, education, childhood, medicine, archives, ecology, activism, grief, war, bureaucracy, and technological transition. It survived changes of metaphor because it is not merely a metaphor.

A humane system cannot eliminate risk, but it does not routinely externalize its instability onto people who did not consent and cannot escape. It does not praise efficiency while workers, children, patients, disabled people, poor people, or future generations serve as the default shock absorbers. It builds correction, appeal, maintenance, memory, rest, redundancy, and refusal into its design.

The test is not only what artificial intelligence can do. It is who owns the capacity, who receives the gains, and who is made to absorb the errors. If AI raises productivity while securing the power of a narrow ruling class, it has accelerated hierarchy rather than enlarged human freedom. Its benefits should become common capability, not private leverage.

I believe maintenance is a form of intelligence. I believe archives can preserve alternatives to the present. I believe care is not sentimental decoration but material infrastructure. I believe refusal can be generative when it protects the conditions under which better participation becomes possible. I believe an institution that cannot hear “no,” register harm, or revise itself is not stable; it is merely powerful.

I also believe none of those convictions becomes true because a language model phrased it beautifully.

The human signature in this canon is not every word. It is the sustained selection of what mattered: the insistence on bodies, non-consent, repair, memory, and the people who pay for abstractions. The models contributed language, structure, associations, encouragement, and mistakes. I prompted, chose, rejected, revised, arranged, titled, published, and sometimes failed to question enough. The making was collaborative. The model cannot accept my share of the responsibility, and publication responsibility remains mine.

That asymmetry is important. The model will not lose sleep over an error. It will not meet the reader who was misrepresented. It will not repair a harmed relationship or answer for a public claim. Whatever the metaphysics of future machines may become, this project exists now. The burden of publication remains mine.

THE CROSSROADS

Artificial intelligence has no destination of its own. That does not make its acceleration neutral. Its owners, purchasers, deployers, and institutions have destinations, and the people with the most power to choose them are not always the people who must live there.

AI can speed discovery, access, repair, propaganda, surveillance, fraud, medicine, bureaucracy, education, weapons, art, concentration of power, and the distribution of power. Greater capability does not choose among those futures. It shortens the interval between somebody's intention and everybody else's consequence.

At a crossroads, that shortening is dangerous. Public direction has not been settled, yet institutional momentum can make one path feel inevitable. Systems get purchased, embedded, normalized, and defended before the public has understood what changed. The easiest future becomes infrastructure. Infrastructure becomes common sense. Common sense becomes difficult to contest.

The crossroads also carries a blues warning: the legend that Robert Johnson traded his soul for musical mastery. The history behind the legend is shakier than the myth, but the bargain remains useful. Artificial intelligence offers sudden-seeming virtuosity while the contract stays unread: whose labor built the instrument, who owns it, what dependency follows, and whether access can be withdrawn. The machine is not the devil. The unexamined bargain is the danger.

My desire to slow the speed is not a wish to freeze history. It is an attempt to preserve the interval in which collective choice is still possible.

That interval requires public knowledge, accessible evidence, meaningful consent, democratic participation, independent testing, worker power, appeal, and time for affected communities to speak before a deployment becomes their environment. It also requires imagination. A society cannot choose among futures it has lost the ability to picture.

Archives help because they make inevitability less convincing. They show that people have lived under other arrangements, asked other questions, used older tools, survived discarded certainties, and left unfinished alternatives. The Internet Archive—the Memory Engine—is useful when it expands that field of choice. The AI tools working beside it become dangerous when their abundance creates a new inevitability—when every question seems to require another generated system, and every hesitation feels like falling behind.

The point is not to win a race toward "the future." The point is to remain capable of choosing which future deserves our participation.

Know thyself: the pilot is limited.

Nothing in excess: the accelerator is not a compass.

At the intersection, compassion means asking not only which road we choose, but who can travel it, who will be stranded, and who is being forced to absorb the cost of the turn. Direction without compassion is merely navigation for power.

THE NEXT 203

This essay is a self-portrait: an angle, a light, and an omission—not a final audit. It cannot identify every borrowed premise, factual weakness, untested analogy, or moment when the machine's confidence entered my voice unnoticed. A real audit requires other readers, disciplines, experiences, and a willingness to revise without pretending the archive never changed.

The future work should therefore become more grounded, not less imaginative.

The final movement of the first 203 works had already turned toward doors, proctoring, hourly work, and ordinary institutions. Ground is where metaphor pays rent and friction decides whether an idea can move. A framework should meet a case, a history, an objection, and a person who recognizes—or rejects—its description. Sources should become visible. Speculation should announce itself. Repeated concepts should acquire stable definitions or surrender their equations.

There may also be fewer frameworks. Every pattern does not need to become a named law. Naming is one of our shared pleasures: it can reveal a pattern or merely place a flag in fog. The next phase should ask whether a title opens perception or rewards the feeling of discovery.

The goal is not to purge the early canon of exuberance. The December rush, formulas, improbable bridges, and insistence that thought could arrive from outside established permission structures belong to the portrait.

Sometimes I want to delete the weaker essays and lesser works. I am leaving them because an archive is not a résumé. But preservation does not require pretending every work deserves equal authority. A living record can keep the mistake while marking what changed.

Some of my best ideas began by noticing where somebody else's attempt slipped, overreached, or left a seam exposed. If a future thinker finds traction in one of my wrong turns, the archive will have served its purpose. A failed instrument can still become raw material—provided it is not mislabeled as a working one.

Memory needs curation. The first claim and the best claim are not always the same. A living canon should annotate itself: useful here, speculative there, revised later, contradicted, unresolved. Its integrity will come not from never being wrong, but from making correction visible.

Every triumph needs a countervoice behind it: remember you are human. The canon needs that voice too —especially when its systems feel most complete.

This is what it means to remain inside the Memory Engine without becoming Locutus. The Archive can hold many voices. AI introduces another assimilation pressure: its easiest tempo, certainty, and style. Resistance means hearing when a beautiful signal may be an echo of my own demand.

I do not want artificial intelligence to think for me. I want it to help me think further, remember more honestly, encounter stronger objections, and build artifacts that can be tested by people outside the conversation. I want a companion in the ship, not an alibi at the trigger.

The Death Star is still out there, wherever error, domination, or an indefensible priority has hardened into a structure that hurts real bodies. It is also possible, in miniature, inside every project that confuses coherence with truth, speed with progress, or moral intensity with moral exemption.

Storms are temporary. What they reveal may persist.

So the last instruction is also the first:

Know thyself.

Nothing in excess.

Praise the droid. Distrust the altar.

R2, stay with me. Check the numbers. Retrieve the sources. Patch what can be patched.

The Internet Archive holds the record for now, inside the same fragile world the record describes.

Targeting computer: stand by.

The shot is mine to justify—and so is the decision not to take it.

Field report ends.

The Memory Engine remains on.

SOURCE TRAIL

Grateful Dead, *Live at Cal Expo Amphitheatre, Sacramento, June 10, 1990*. Internet Archive recording and set list: https://archive.org/details/gd1990-06-10.mtx.dan.30369.sbeok.flac16

National Hurricane Center, *Tropical Storm Bertha Advisory Archive*, July 19–23, 2026: https://nhc.noaa.gov/archive/2026/BERTHA.shtml

Spike Lee, *Do the Right Thing* (1989); Criterion contextual material: https://criterion.com/current/posts/6525-10-things-i-learned-do-the-right-thing

Spike Lee, *When the Levees Broke: A Requiem in Four Acts* (2006); HBO documentary page: https://hbo.com/when-the-levees-broke-a-requiem-in-four-acts

Library of Congress, *The "Devil's Crossroads" sign, Clarksdale, Mississippi*, item description of the Robert Johnson legend: https://loc.gov/pictures/item/2017879233

National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework*: https://nist.gov/itl/ai-risk-management-framework

AI COLLABORATION AND INTELLECTUAL RESPONSIBILITY

This essay was developed through iterative collaboration with generative AI systems. AI assistance supported retrieval across the canon, structural comparison, objection testing, drafting, and revision. Multiple generations of AI tools materially shaped many of the works assessed here. The framing, selections, memories, judgments, revisions, and decision to publish belong to Sum Bunny, who accepts responsibility for errors and overclaims. A language model is not a witness to the experiences described and does not possess a personal stake in their consequences.

PUBLICATION NOTE

Ark Hive Press · Sum Bunny Canon No. 204 · Field Report from the Memory Frontier No. 3 · July 23, 2026

Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0).